Exhibit M

IN THE UNITED STATES DISTRICT COURT
FOR THE EASTERN DISTRICT OF TEXAS
SHERMAN DIVISION

INTELLECTUAL VENTURES I LLC and INTELLECTUAL VENTURES II LLC,	)	
Plaintiffs,	)	**C.A. No. 4:24-cv-00980-ALM**
v.	)	
	)	**JURY TRIAL DEMANDED**
AMERICAN AIRLINES, INC.	)	
Defendant.	)	

REBUTTAL DECLARATION OF EYAL DE LARA

I, Eyal de Lara, declare as follows:

1. My name is Eyal de Lara. I am the Chair of the Department of Computer Science at the University of Toronto where I have also been a professor since 2002.

2. If called to testify under oath in court, I could and would testify competently to the facts stated herein.

3. I have been retained by Intellectual Ventures I LLC and Intellectual Ventures II LLC ("IV") as an independent expert consultant in this proceeding. I have been asked to submit this declaration on behalf of IV.

4. I have been asked to provide rebuttal opinions regarding U.S. Patent Nos. 7,257,582 ("'582 Patent"), 7,949,785 ("'785 Patent"), 8,332,844 ("'844 Patent"), and 8,407,722 ("'722 Patent") (together, "Asserted Patents"), in response to the declaration of Michael T. Goodrich, Ph.D. My opinions are set forth below.

5. I am being compensated at a rate of $600 per hour for my work in this proceeding. My compensation is not contingent on the nature of my analyses and opinions, the presentation of my findings in testimony, or the outcome of this proceeding.

6. All of my opinions stated in this Declaration are based on my own personal knowledge and professional judgment. I am over 18 years of age and, if I am called upon to do so, I would be competent to testify as to the matters set forth in this declaration.

I. QUALIFICATIONS

7. I am currently the Chair of the Department of Computer Science at the University of Toronto, a position I assumed in 2022. The University of Toronto's Department of Computer Science is consistently ranked among the top departments in the world, and as Chair, I am responsible for the academic, administrative, and strategic leadership of the department. This includes oversight of faculty recruitment, curriculum development, research initiatives, and graduate and undergraduate programs.

8. In 2002, I was awarded a Ph.D. in Electrical and Computer Engineering from Rice University in Houston, Texas. My doctoral dissertation was entitled "Component-based Adaptation for Mobile Computing." This work addressed the design of software systems capable of adapting to the changing conditions of mobile and pervasive computing environments, a research area that, at the time, was emerging as central to the development of modern computing infrastructure. As described in greater detail below, my professional career has been directed toward experimental research distributed systems with a focus in cloud, mobile and pervasive computing systems.

9. Prior to my doctoral studies, I earned a Bachelor of Science degree in Computer Science in 1995 from the Instituto Tecnológico y de Estudios Superiores de Monterrey (ITESM), in Mexico. ITESM is widely regarded as one of the leading technical universities in Latin America. Following my undergraduate degree, I pursued graduate study at Rice University, where I obtained a Master of Science degree in Electrical and Computer Engineering in 1999. My master's thesis

concerned distributed computing systems, with a specific focus on the use of shared memory in page-based software environments.

10. During the course of my graduate studies, I also gained experience outside of academia. In 1998, I was selected for a competitive summer research internship at Xerox Palo Alto Research Center (PARC), located in Palo Alto, California. Xerox PARC has historically been one of the most influential industrial research centers in the history of computer science and engineering, having pioneered innovations such as the graphical user interface and Ethernet.

11. After completion of my Ph.D. in 2002, I was appointed as an Assistant Professor in the Department of Computer Science at the University of Toronto. I was subsequently promoted to the rank of Associate Professor in 2007, full Professor in 2015, and Chair of the Department in 2022. I also hold a cross-appointment in the Department of Electrical and Computer Engineering at the University of Toronto, reflecting the interdisciplinary nature of my research and teaching.

12. My primary research areas are cloud computing and mobile computing. Within these areas, I have made contributions to a wide range of topics, including system virtualization, edge computing, application scaling, indoor localization, mobile security, and continuous sensing using mobile devices. Cloud computing and edge computing are concerned with the placement and distribution of computing and storage resources in networked environments, typically Internet-based, and have become fundamental to modern software services and infrastructures. My research has spanned from detailed technical investigations, such as the allocation of graphics processing units (GPUs) in virtual machine environments, to broader architectural questions relating to the provision of software and data through remote and distributed computing systems.

13. In addition to my research contributions, I have been active in the leading international professional societies in computer science, namely the Association for Computing

Machinery (ACM) and the Institute of Electrical and Electronics Engineers (IEEE). These organizations are the two principal bodies governing the professional and scholarly activities of computer scientists and electrical engineers worldwide. My involvement has included publication of research, leadership roles, and recognition through awards. For example, in 2024, I co-authored a paper that received the Best Paper Award at the IEEE/ACM Symposium on Edge Computing. I was elected a Senior Member of the IEEE in 2019, an honor reserved for those with significant professional accomplishments. In 2020, I was the recipient of the ACM SIGMOBILE Distinguished Service Award, which is granted by ACM's Special Interest Group on Mobile Computing to individuals who have provided exceptional service to the community.

14. My research and scholarly contributions have been recognized by multiple awards and honors. These include the ACM EuroSys Test of Time Award in 2019, which recognizes work of enduring influence in the field of computer systems, and the CACS/AIC Outstanding Young Computer Science Researcher Prize in 2012, which recognizes early-career excellence. I have also received an NSERC Discovery Accelerator Supplement, faculty research awards from IBM and VMware, four Best Paper Awards, and three Best Paper Honorable Mentions. At the University of Toronto, I was awarded the Connaught Innovation Award in 2020 and 2024 and was designated one of the University's Inventors of the Year in 2012.

15. I am also the author of a monograph on location technologies, published as part of the Morgan and Claypool Synthesis Lecture series on Mobile and Pervasive Computing. From 2014 to 2019, I served as Editor-in-Chief of GetMobile, the flagship publication of ACM SIGMOBILE. In this role, I was responsible for overseeing the publication of scholarly articles and surveys on mobile computing, guiding editorial direction, and ensuring quality and relevance to the field. In addition, I have published extensively in peer-reviewed venues on cloud, mobile,

and edge computing, and I have served on the editorial boards of multiple journals published by IEEE and ACM in these subject areas.

16. Beyond my research, I have extensive experience in teaching and curriculum development, particularly in the area of cloud computing. In 2012, I created and introduced the graduate-level course ECE1779: Introduction to Cloud Computing as part of the Master of Engineering program offered by the Department of Electrical and Computer Engineering at the University of Toronto. This course provides graduate students with both theoretical and practical foundations in cloud computing. The syllabus includes coverage of enabling technologies such as, virtualization, scalability, fault tolerance, security, load-balancing, 3-tier architectures, and distributed processing paradigms, such as batch and stream processing. It combines lectures with practical programming assignments using commercial cloud platforms, such as Amazon Web Services' EC2 and Lambda. Since its inception, I have taught 13 offerings of the course, expanding it to two offerings per year beginning in 2019 to meet increasing demand. The syllabus has been updated over time to reflect changes in the technology. The most recent version of the course, taught in Spring 2022, focused on Amazon technologies including EC2, S3, Lambda, DynamoDB, RDS, CloudWatch, CloudFront, API Gateway, and Elastic MapReduce, and Apache Storm.

II. MATERIALS REVIEWED

17. In forming the opinions provided in this Declaration, I have reviewed the Declaration of Dr. Goodrich, including documents cited in that declaration that relate to the patents for which I am offering opinions on. The opinions contained in this declaration are based on the documents I reviewed, as well as my professional experience, education, and knowledge. In forming my opinions expressed in this declaration, I have reviewed the '582 Patent, the '785 Patent, the '844 Patent, and the '722 Patent, as well as their prosecution histories. I have also reviewed other materials referred to in this declaration in support of my opinions.

18. My opinions contained in this Declaration are based on the documents I reviewed and my knowledge and professional judgment. My opinions have also been guided by my appreciation of how a person of ordinary skill in the art would have understood the state of the art, the prior art, and the claims and the specification of the '582 Patent, the '785 Patent, the '844 Patent, and the '722 Patent at the time of the invention.

III. LEGAL STANDARDS

19. I have been informed of legal standards by counsel for IV in forming the opinions provided in this Declaration.

20. I am not an attorney and I do not claim to be one. Counsel for IV has explained to me legal principles to keep in mind in forming my opinions provided in this declaration. Those legal principles are provided below.

21. I understand that claim construction is solely a matter for a court to decide and, in general, the customary and ordinary meaning of the claim terms used in the patent to one of ordinary skill in the art is determined in the context of the patent's specification and the file history.

22. I understand that a person of ordinary skill in the art is a hypothetical person who is presumed to have known the relevant art at the time of the invention.

23. I understand that factors that may be considered in determining the level of ordinary skill in the art may include: (1) type of problems encountered in the art; (2) prior art solutions to those problems; (3) how quickly innovations are made; (4) sophistication of the technology; and (5) educational level of active workers in the field. In a given case, every factor may not be present, and one or more factors may predominate.

24. I understand that a person of ordinary skill in the art is a person of ordinary creativity, and would have the capability of understanding the scientific principles applicable to the pertinent art.

25. I understand that claims are construed from the perspective of a person of ordinary skill as of the effective filing date of the patent application.

26. I understand that persons of ordinary skill in the art are deemed to read the claims in the context of the entire patent, including the specification and prosecution history. In other words, the terms are not considered in a vacuum.

27. I understand that, in the context of claim construction, the specification is the single best guide to the meaning of the claim terms.

28. I understand that claim terms should be given their ordinary and customary meaning within the context of the patent in which the terms are used, *i.e.*, the meaning that the term would have a person of ordinary skill in the art in question at the time of the invention in light of what the patent teaches.

29. I understand that the plain and ordinary meaning is determined from the language of the claims, the specification, and the prosecution history of the patent at issue.

30. I understand that, in construing a claim term, one looks primarily to the intrinsic patent evidence, which includes the patent abstract, specification, claims, and figures, and its prosecution history.

31. I understand that extrinsic evidence may also be useful in interpreting patent claims when the intrinsic evidence itself is insufficient.

32. I understand that the usual and customary meaning of a claim term can be altered by the patent applicant if they choose to act as their own "lexicographer" and clearly set forth in the patent a different meaning of a claim term.

33. I understand that the meaning of a claim term can also be altered by a clear and unequivocal disavowal by the patent applicant made regarding the meaning or scope of the term.

34. I understand that if a claim term is ambiguous or unclear, the term must be construed to determine how a person of ordinary skill in the art would have resolved in light of the rest of the patent specification, patent claims, and the application's prosecution history.

35. I understand that a claim is not indefinite, as long as it, having been read in light of the intrinsic evidence, informs one of skill in the art at the time of the invention about the scope of the invention with reasonable certainty.

36. I understand that it is improper to import limitations from embodiments in the specification.

37. I understand that it is improper to import limitations from other parts of the claims and rendering a claim term duplicative.

38. I understand that it is improper to import additional or different language from other independent claims that would render such claims superfluous.

IV. PERSON OF ORDINARY SKILL IN THE ART

39. In his declaration, Dr. Goodrich offers the following opinion as to the standard of a person of ordinary skill in the art for each of the '582 Patent, the '785 Patent, the '844 Patent, and the '722 Patent.

> Based on my review of the Patents-in-Suit and my consideration of the abovementioned factors, it is my opinion that a person of ordinary skill in the art ("POSITA") as of the respective priority dates for the Patents-in-Suit would have had a Bachelor's degree in electrical engineering, computer science, or the like, and/or two or more years of industry experience in networking and distributed computing. Additional experience may make up for a lack of education and vice versa. For example, each of the Patents-in-Suit discloses inventions involving computational devices and networking.

Goodrich Decl. at ¶ 21.

40. For purposes of this declaration, I have applied this standard offered by Dr. Goodrich. Under Dr. Goodrich's definition, I am a POSITA for each of the '582 Patent, the '785 Patent, the '844 Patent, and the '722 Patent.

V. '582 PATENT

A. Background of the '582 Patent

41. The '582 Patent is directed to "the field of sharing data and workload between possibly heterogeneous computer systems." It explains that "more specifically, it deals with a way to split the performance of a given task among a plurality of processing units which can all access, directly or indirectly, the input data and the devices on which the output data is to be stored." The patent further notes that "[s]ort applications, statistical analysis batch applications and report writing applications and database queries are examples of applications that can readily enjoy this invention." '582 Patent at 1:11-21.

42. In relation to this subject matter, my research has addressed cluster computing, parallel computing, data-intensive computing, distributed computing, and computing on heterogeneous devices. I highlight below my publications most directly related, using the indexing convention from my CV where C, J, P, and W correspond to conference, journal, patent, and workshop publications, respectively:

43. C31, C32, C35, J11, P5: Introduce mechanisms for fast virtual machine cloning and apply this capability to parallel computation on clusters of servers, as well as to run web and database servers that can be dynamically resized to match client demand.

44. C63, W33: Investigate shared-memory frameworks for running parallel high performance scientific workloads on clusters of servers.

45. J12: Leverages distributed computing techniques to execute interactive queries over large datasets. C29: Proposes a framework for running data-intensive applications on

virtualized datacenters. This framework enables efficient file sharing across virtual machines (VMs) on the same physical host and optimizes VM placement within the cluster to maximize sharing opportunities.

46. C17, C21, J3: Introduce new primitives and frameworks for executing applications in geo-distributed computing environments.

47. C24: Presents a programming framework for applications that execute on collections of heterogeneous devices.

B. Disputed Terms

48. For the '582 Patent, I understand that Dr. Goodrich offers opinions as to the meaning of the following terms to a POSITA: (1) "partition"; (2) "description of all of said partitions"; (3) "on a first-come/first-served basis"; and (4) "simultaneously executing at least a respective one of the subtasks of the computer-executable process in each of at least some of said processors on a respective one of the partitions."

49. I have reviewed the opinions offered by Dr. Goodrich in paragraphs 60-82 of his declaration relating to the '582 Patent disputed claim terms. For the reasons below, I disagree with Dr. Goodrich's opinions as to the meaning of those terms.

Claim Term	***IV's Proposed Construction***	***American's Proposed Construction***
partition (claim 1)	Plain and ordinary meaning, no construction necessary	well-defined part of the input file

50. I understand that Dr. Goodrich alleges that the POSITA would have understood that the patentee defined "partition" in the specification as "a well-defined part of the input or output."

51. Dr. Goodrich based this conclusion on the '582 Patent disclosure:

> A logical partition in this context is a well-defined part of the input or output. A very simple way to define a partition, which in many cases is the most efficient, would be to define the partitions as consecutive ranges on the input or output, ranging from one relative byte address to another relative byte address or from one relative track address to another.

'582 Patent at 3:36-41.

52. However, the same paragraph also discloses:

> Other partition definitions can be used with no impact on the rest of the embodiment.

'582 Patent at 3:46-47.

53. This disclosure indicates that the patent does not place restrictions on the nature of the partition leading the POSITA to understand the term in its plain an ordinary meaning.

Claim Term	***IV's Proposed Construction***	***American's Proposed Construction***
descriptions of all of said partitions	Plain and ordinary meaning, no construction necessary	statements giving a characteristic(s) of all of the well-defined parts of the input file for use in distributing the load without a special load process, wherein such statements are distinct from the input file itself

54. Dr. Goodrich claims "descriptions of all of said partitions" would have been understood by a POSITA when reading the claims in view of the specification and the prosecution history as "statements giving a characteristic(s) of all of the well-defined parts of the input file for use in distributing the load without a special load process, wherein such statements are distinct from the input file itself."

55. The declaration offers the following opinion at paragraph 64:

> 64. Second, during prosecution of the '582 Patent, the applicant responded to the examiner's rejection of the pending claims by distinguishing prior art and stating that "The amended claims, refer to a distribution of a description of the work to be done. The sharing process can use such a description to distribute the load without a special load process." First Amendment, March 2, 2007, at 7. It is

> my understanding that a patent applicant's statements during prosecution to distinguish the claims over the prior art may show that the applicant disavowed claim scope.

56. The text surrounding this cited quote in the '582 Patent file history says:

> The primary difference between the instant invention and the processes disclosed in US 5,603,028 of Kitsuregawa and 5,357,632 of Pian is that these systems rely on a special control process that uses load information to distribute the load between processors that share the load. With the instant invention as defined in the claims there is no such special process. The prior art's load information is not created with the process of the instant invention. Instead, the load sharing is done as a byproduct of the fact that the load-sharing process take parts of the load on a first-come/first-served basis.
>
> A comparison would be to a road intersection where, according to the prior art, there is a traffic light that determines who can go when. The instant invention is more like such an intersection with a four-way stop so that the individual drivers determine who can go and when.
>
> This is a major improvement since in addition to eliminating the control process it also eliminates the need to collect and maintain load information, which it is very difficult to do and almost impossible to define so as to anticipate all possible processors that might execute the subtasks.
>
> The amended claims, refer to a distribution of a description of the work to be done. The sharing process can use such a description to distribute the load without a special load process.

57. The text removes the requirement for a centralized load process, or more precisely the existence of a centralized scheduling process that uses load information to make scheduling decisions. The amendment, however, does not preclude centralized coordination or the use of load information to make scheduling decisions; it simply makes this an option as opposed to a hard requirement. This is clear in the amendment which states the "*the sharing process CAN use such a description to distribute the load without a special load process.*"

58. This disclosure indicates that the patent does not place restrictions on the nature of the term, leading a POSITA to understand it in its plain an ordinary meaning.

Claim Term	*IV's Proposed Construction*	*American's Proposed Construction*
on a first-come/first-served basis	Plain and ordinary meaning, no construction necessary	selecting the earliest unprocessed partition for execution without the use of a control process that uses load information for such selection

59. I understand that Dr. Goodrich alleges that the *POSITA would have understood "first-come/first-served basis" to mean "selecting the earliest unprocessed partition for execution* ***without the use of a control process*** *that uses load information for such selection."*

60. Dr. Goodrich opines in paragraphs 75 and 76:

> A POSITA would have also understood that whether a partition is the "earliest" unprocessed partition would have been determined by the ordering of the partitions determined in step (a) of claim 1.
>
> During prosecution, the applicant distinguished prior art, stating that "[the prior art] systems rely on a special control process that uses load information to distribute the load between processors that share the load. With the instant invention as defined in the claims there is no such special process. . . Instead, the load sharing is done as a byproduct of the fact that the load-sharing process take parts of the load on a first-come/first-served basis." First Amendment, March 2, 2007, at 6. From the applicant's distinction of the claims over the prior art, a POSITA would have understood that this selection is done without the use of a control process that uses load information for such a selection and instead is done according to an ordering of the partitions. determined in step (a) of claim 1 ("automatically determining file allocation and logically subdividing records of said input file into a plurality of partitions").

61. It is my opinion that there is no evidence provided in the patent that the "selecting the earliest unprocessed partition" must be determined by the information shared in step (a). The "four-way stop" example provided in the file history is illustrative of an approach where processes make decisions based on real time information. Moreover, the patent description allows for the inclusion of an (external) scheduler that considers dependencies:

> The various parts of the Equivalent Process 101 are depicted here as consecutive steps in one procedure but alternative embodiments could replace them by dependent 15 tasks controlled by a job scheduler. In this case, the wait step would be replaced by a dependency of the Merge step 105 on the completion of Sub Task 103, Sub Task 201 and Sub Task 202.

'582 Patent at 3:13-19.

> When the split step terminates, the various Sub Tasks (in this case: 103, 202.203) can be activated. This activation can 65 be initiated by the split step 102 itself or by an external scheduler.

'582 Patent at 3:64-67.

62. In this light, it is my opinion that the POSITA would have understood meaning of the term to have its plain an ordinary meaning.

Claim Term	***IV's Proposed Construction***	***American's Proposed Construction***
simultaneously executing at least a respective one of the subtasks of the computer-executable process in each of at least some of said processors on a respective one of the partitions	Plain and ordinary meaning, no construction necessary Not indefinite	Indefinite

63. I understand that Dr. Goodrich alleges that the disputed claim term "simultaneously executing at least a respective one of the subtasks of the computer-executable process in each of at least some of said processors on a respective one of the partitions" is indefinite because a POSITA could not have ascertained its meaning with "reasonable certainty." I disagree.

64. Claim 1 recites in part a method "of effecting on a preexisting input file a computer-executable process" that includes "a *plurality of subtasks*." '582 Patent at 6:44-46 (emphasis mine). It also recites "distributing descriptions of all of said partitions to each of *a plurality of subtask*

processors." '582 Patent at 6:50-51 (emphasis mine). Claim 1 further recites "simultaneously executing *at least a respective one of the subtasks* of the computer-executable process *in each of at least some of said processors* on a respective one of the partitions." '582 Patent at 6:52-55.

65. I understand that a "plurality" generally means two or more. Thus, as recited, claim 1 recites two or more subtasks and two or more subtask processors.

66. Dr. Goodrich opines that the "simultaneously" limitation recites executing only one subtask and is thus indefinite because this limitation is inconsistent with the '582 Patent specification that describes parallel processing. Goodrich Decl. at ¶ 82. However, in my opinion, a POSITA would understand that claim 1 requires at least two or more subtasks to be executed "simultaneously." This is because at least one subtask must be executed in *each* of the subtask processors. The claim requires at least two subtask processors, and in order to satisfy this limitation, there must be at least two subtasks, one for each subtask processors. This conclusion is entirely consistent with the '582 Patent's specification regarding parallel processing.

VI. '785 PATENT

A. Background of the '785 Patent

67. The '785 Patent concerns virtual network technology "that allows local and remote entities to communicate and collaborate from various locations." '785 Patent at 5:48–51. The '785 Patent provides a basic overview of a Virtual Communication Network or "VCN" in Figure 4. A VCN allows devices, such as devices A, B, and X, to communicate as if they were on the same physical local network. As shown in Figure 4, devices A, B, and X may direct connections to each other through an application IP interface ("IPa," "IPb," and "IPx") within a virtual domain, where "dashed lines represent direct communication paths seen to applications running on A, B and X." *Id.* at 9:36-41. Devices A and B may have dynamic or static private network addresses, and device X may be coupled directly to the Internet and may have a public network address. *Id.* at 9:15-17.

68. My related research includes work on wireless networking, virtual networking, and secure communications:

69. C43, C49, W9: Developed new multi-hop ad hoc network architectures leveraging node mobility to improve connectivity. C53, W30, W31: Extended routing and MAC protocols in mobile ad hoc networks to enhance service discovery, reduce self-interference, and improve flow control.

70. C25, C30, J10, P6: Combined virtualization technology with specialized hardware to enable servers to maintain network presence while reducing energy consumption. P5, J11, C31, C32, C35: Leveraged network interface virtualization to preserve communication during virtual machine migration and cloning in datacenter environments.

71. C34, C41, J13, W25: enable secure communication between parties by exploiting characteristics of the radio environment to prevent eavesdropping and impersonation attacks, such as man-in-the-middle.

B. Disputed Terms

72. For the '785 Patent, I understand that Dr. Goodrich offers opinions as to the meaning of two claim terms: (1) "network address"; and (2) "network route director."

73. I have reviewed the opinions offered by Dr. Goodrich in paragraphs 36-50 of his declaration relating to the '582 Patent disputed claim terms. For the reasons below, I disagree with Dr. Goodrich's opinions as to the meaning of those terms.

Claim Term	*IV's Proposed Construction*	*American's Proposed Construction*
network address	Plain and ordinary meaning, no construction necessary	Internet protocol or IP address

74. In my opinion, a POSITA would not equate a network address with an IP address. The dictionary definition included in Dr. Goodrich's declaration at paragraph 44 confirms this

understanding when it offers a definition of a network address as the "network *portion*" of an IP address, not the IP address itself. Thus, any interpretation that a network address is identical to an IP address is incorrect, and a POSITA would have understood this. I reviewed the disclosures cited by Dr. Goodrich in this declaration, and while some embodiments may use network address interchangeably with an IP address, there is no definition in the specification that defines the two terms to have identical meaning.

75. Further, claim 30 recites different types of network addresses, including "public" network addresses, "private" network addresses, and "virtual" network addresses. I understand that Dr. Goodrich only offers opinions on the single recitation of "network address" at 36:52, not these other recitations of a type of network address. A POSITA would appreciate that the "network address" recited at 36:52 would be broader than the particular types of network addresses recited elsewhere in claim 30.

76. Dr. Goodrich's reliance on a reexamination history at paragraphs 42 and 43 is not dispositive, and further it is unclear which address his citation is even discussing because claim 30 recites that the DNS server returns a network address, a private network address, and a virtual network address. While there are embodiments that perhaps disclose returning an IP address for one or more of those addresses, there is nothing in the claim or specification that would necessarily limit them to being only IP addresses.

Claim Term	***IV's Proposed Construction***	***American's Proposed Construction***
network route director	Plain and ordinary meaning, no construction necessary	a publicly addressable device configured to route encapsulated packets to and from entities located in a private network portion of a virtual network

77. I understand that Dr. Goodrich opines in paragraph 46 of his declaration that this term should be construed to mean "a publicly addressable device configured to route encapsulated packets to and from entities located in a private network portion of a virtual network." I disagree.

78. Claim 30 recites in part a DNS server for the recited virtual network that is "configured to receive a DNS request from a first device in the virtual network" and "*return a network address associated with a network route director.*" '785 Patent at 36:50-53 (emphasis mine). Thus, Claim 30 recites a DNS server that receive a DNS request and "return[s]" a "network address," where the network address is "associated with a network route director." The network route director limitation is not recited elsewhere in claim 30. Thus, A POSITA would understand that the only requirement of claim 30 is that the network address must be "associated" with the network route director.

79. In my opinion, Dr. Goodrich's opinion is wrong because it reads limitations into claim 30 that are not supported by intrinsic evidence. For example, claim 39 depends on claim 38, and recites that "the route director is a *public* network route director." '785 Patent at 37:39-40. This language would indicate to a POSITA for example that a network route director does not have to be publicly addressable.

80. Dr. Goodrich's opinion is further incorrect as to the language that reads "configured to route encapsulated packets to and from entities located in a private network portion of a virtual network" because there is nothing in the claims that require that the network route director be configured in such a manner. Dr. Goodrich's proposed construction effectively re-writes claim 30 to require the network route director to route messages to the recited second device using the private network address and virtual network address returned by the DNS server, even though such a requirement is recited nowhere in claim 30. Dr. Goodrich cites for example a disclosure at

column 13 lines 51 to 59, but a POSITA would appreciate that this is a single embodiment described in the '785 Patent specification, and that the passage in question states that a Network Route Director "facilitate[s]" routing traffic, which is different than routing encapsulated packets under Dr. Goodrich's proposed construction.

VII. '844 PATENT

A. Background of the '844 Patent

81. The '844 Patent is directed to technology for "root image caching and indexing for block-level distributed application management." The technology involves storing blocks of a root image on a first storage unit and storing blocks of leaf images on respective second storage units. The leaf images include additional data blocks not previously contained in the root image as well as changes made by respective compute nodes to the blocks of the root image. '844 Patent at 2:35–43.

82. My related research includes work on distributed file systems, distributed block storage devices, and container provisioning frameworks:

83. W3: Introduced PathFS, a new distributed file system designed for hierarchical edge deployments. PathFS adopts a tree-like structure with cloud datacenters at the root, edge datacenters as leaves, and intermediate network datacenters. To address limited storage at lower levels, PathFS employs on-demand block-level partial replication: files are persisted in the cloud (root) while intermediate and edge layers act as temporary caches.

84. C35, C39, W24: Introduced WanDisk, a peer-to-peer distributed storage layer for efficiently managing VM disk state across WAN links. WanDisk uses block-level replication (splitting virtual disks into versioned chunks replicated across peers) and copy-on-write (updating chunks locally and deferring synchronization). These mechanisms allow VMs to migrate quickly across WANs while ensuring consistent, on-demand disk state.

85. C11: Introduced Starlight, an accelerator for container provisioning. Starlight decouples provisioning from development by redesigning the deployment protocol, filesystem, and image storage format. While maintaining the familiar layered structure of container images, Starlight employs a more efficient representation for network deployment, eliminating duplicate data and reducing network roundtrips. Its filesystem provides containers with a merged view combining multiple read-only layers and a single read-write layer.

B. Disputed Terms

86. I understand that, for the '844 Patent, Dr. Goodrich offers opinions as to only a single claim term, "root image."

87. I have reviewed the opinions offered by Dr. Goodrich in paragraphs 51-59 of his declaration relating to the sole '844 Patent disputed claim term. For the reasons below, I disagree with Dr. Goodrich's opinion.

Claim Term	***IV's Proposed Construction***	***American's Proposed Construction***
root image	Plain and ordinary meaning, no construction necessary	a read-only base set of data blocks, operating beneath the file system, that provide the common portion of the application environment

88. Claim 7 recites, in part, "*storing blocks of a root image* of said compute nodes on a first storage unit" and "storing leaf images" that "*includ[e]* only additional data blocks not previously contained in said root image and *changes made by respective compute nodes to the blocks of the root image*." '844 Patent at 11:29-30 (emphasis mine). A root image includes blocks of data, and changes to those blocks are stored in the leaf image.

89. The portion of Dr. Goodrich's interpretation that includes "read-only" language is incorrect because it reads out embodiments of the '844 Patent that disclose root images that are not read-only. For example, Figure 2 shows a system **200** that includes a first storage unit **240** that

"store[s] blocks of a root image of an application environment." '844 Patent at 5:27-28. The specification describes that, "*in embodiments where the first storage unit **240** is read-only*, the root image will not have to be re-indexed because the contents of the root image do not change." '844 Patent at 7:34-37. In other words, the specification describes that not all embodiments include a "read-only" root image.

90. The next portion of Dr. Goodrich's interpretation, "operating beneath the file system," is not based on any claim language. The proposed language thus appears to be imported from embodiments of the '844 Patent, which Dr. Goodrich cites for example the disclosures at 7:58-62 and 5:51-58. However, these disclosures in my opinion relate to embodiments of the '844 Patent. In particular, the description at 7:58-62 describes that this embodiment is *able* to operate at a layer beneath the file system, not that it must do so. The disclosure at 5:51-58 is consistent with this understanding.

91. The portion of Dr. Goodrich's interpretation that reads "that provide the common portion of the application environment" should be rejected because neither the claims nor the specification recite this language. Further, to the extent Dr. Goodrich is correct that lexicography applies here (which it does not in my opinion), Dr. Goodrich's interpretation should be "a read-only base set of data blocks of the application environment."

VIII. '722 PATENT

A. Background of the '722 Patent

92. The '722 Patent is directed to "transferring information through digital networks and in particular to transferring information for remotely updating content at client devices through the digital networks." '722 Patent at 1:24–27. The '722 Patent discloses that a client device receives a live object from an input source, and in response identifies object IDs associated with the object and registers the object IDs with a routing network adapted to send update messages to

nodes in the network. *Id*. at Abstract. Messages from an input source are assigned to one or more categories that is tracked by a gateway, such that when a gateway receives update messages from input sources and identifies the one or more categories, it routes the messages to the appropriate client(s). This routing "utilizes bandwidth efficiently because the update messages are provided to the clients only when the live objects change." *Id*. at 3:18-20.

93. The portion of my research relevant to this subject matter concerns remote document update, content-based routing, stream processing systems, and publish/subscribe frameworks.

94. C3, C4, C8, W12: Propose a stream processing framework that leverages a network of message brokers to route application tuples between operators. The use of message brokers allows data flows to be restructured at runtime, enabling seamless application reconfiguration.

95. C51: Introduces an architecture enabling mobile clients to browse web pages with frequently changing content. This approach allows users to identify portions of interest on a page and receive updates to those portions in an energy-efficient manner as they change over time.

96. C57: Examines the effect of user mobility on a publish/subscribe network of brokers.

B. <u>Disputed Terms</u>

97. For the '722 Patent, I understand that Dr. Goodrich offers opinions as to the following claim terms: (1) "input source"; and (2) identify a category of the update message based on the input source."

98. I have reviewed paragraphs 83-87 of Dr. Goodrich's declaration that offers opinions relating to the '722 Patent. I disagree with Dr. Goodrich's conclusions, and my reasoning and opinions are provided below.

Claim Term	*IV's Proposed Construction*	*American's Proposed Construction*
input source	Plain and ordinary meaning, no construction necessary	information provider and/or dynamic content provider

99. In my opinion, this term should be construed to have its ordinary and customary meaning, "a source of input."

100. Claim 14 recites an "input source" that provides a "data representation" that includes a "live object" to a "client device." While the '722 Patent specification provides examples of an input source including an "information provider" and a "dynamic content provider," it does not limit an "input source" to those examples. Elements 710C and 710D of Figure 7 for example support my opinion. While an input source *can* be an information provider or a dynamic content provider, the specification describes other potential sources of input.

Claim Term	*IV's Proposed Construction*	*American's Proposed Construction*
identify a category of the update message based on the input source	Plain and ordinary meaning, no construction necessary	identify a category of the update message based on the information provider or dynamic content provider but not on the category/topic of the message content

101. I understand that Dr. Goodrich alleges that the POSITA would have understood "identify a category of the update message based on the input source" to mean "identify a category of the update message based on the information provider or dynamic content provider but not on the category/topic of the message content."

102. The description, however, gives various examples of how message category can be determined:

> There are multiple ways to assign the messages into categories. One way is to assign all messages from a given input source 710 into a certain category. Another way is to explicitly specify the category in the object ID for the message. Yet another way is to utilize a

> hashing function or lookup table to partition messages into categories based on object IDs or other values. For example, in one embodiment a hash function is applied to the object ID to generate an integer between 1 and N, and this integer is the message category.

'722 Patent at 18:52-60.

> In one embodiment, the gateway 724 uses a combination of multiple techniques to determine the message categories, node types, and/or mappings. For example, a lookup table can be used to encode a priori knowledge about categories, types, and/or mappings and a hash table can be used to route messages for which there is no a priori knowledge. Continuing this example, assume that certain messages are assigned to a given category based on a table lookup, while other messages are assigned to categories based on a hashing function. In this example, the gateway 724 looks up the object ID (or other information, such as an input source ID) of an arriving message in a lookup table to determine if it has a specified category. If the object ID is stored in the lookup table, the gateway 724 determines the mappings for the category and routes the messages to the nodes of the appropriate types. If the objectID is not stored in the lookup table, the gateway 724 utilizes a hash function on the objectID (or other information) to determine the message category.

'722 Patent at 19:6-23.

103. I have reviewed paragraph 87 in Dr. Goodrich's declaration and disagree with his opinions based on the file history. During prosecution, the Applicant argued in a response dated February 9, 2018: "The Examiner, for example, on page 3 of the Office Action, relies on the above sections of Chandra to allegedly show 'identifying each category based on the category/topic of the message content.' By this statement, the Examiner appears to agree that Chandra teaches identifying categories **based on category/topic of the message content**" (emphasis in original). In reviewing this passage, a POSITA would appreciate that the Applicant was arguing that Chandra did not disclose identifying a category "based on the input source," as recited in the '722 Patent claims. Nothing in the claims or specification preclude identifying a category based on a category/topic of the message content, as long as the limitation "based on the input source" is satisfied.

I hereby declare under penalty of perjury under the laws of the United States of America that the foregoing is true and correct to the best of my knowledge, information, and belief, formed after reasonable inquiry under the circumstances.

Executed on the 26th day of August, 2025, in Toronto, Canada.

Eyal de Lara